GOVERNMENT OF THE REPUBLIC OF VANUATU

PRIME MINISTER'S OFFICE

CERTVU
DEPARTMENT OF COMMUNICATIONS & DIGITAL TRANSFORMATION

PM B 9108 Port Vila, Vanuatu
Tel: (678) 33380

GOUVERNEMENT DE LA REPUBLIQUE DU VANUATU

BUREAU DU PREMIER MINISTRE

CERTVU

DEPARTMENT DE COMMUNICATION ET DE TRANSFORMATION NUMERIQUE

SPP 9108 Port Vila, Vanuatu
Tel: (678) 33380

4 September 2026

Advisory 252: SciPhi-AI R2R Unauthenticated Stacked SQL Injection (CVE-2026-82526).

Release Date: 3rd September 2026
Impact: CRITICAL
TLP: CLEAR

The Department of Communications and Digital Transformation (DCDT) through CERT Vanuatu (CERTVU), provides the following advisory.

This alert is relevant to Organizations, System/Network administrators, and AI/software engineering teams that deploy or operate SciPhi-AI R2R as a self-hosted retrieval-augmented generation (RAG) platform. This alert is intended to be understood by technical users and systems administrators.

What is it?

CVE-2026-82526 is a critical, unauthenticated SQL injection vulnerability in R2R (SciPhi-AI), an open-source, self-described "production-ready" retrieval system (approximately 7,900 GitHub stars) that powers agentic retrieval-augmented generation (RAG) applications via a RESTful API, with enterprise-oriented features including user management, access controls, and Docker/Kubernetes deployment. The flaw is in R2R's vector index creation endpoint (POST /v3/indices): the index_name parameter supplied by the caller is interpolated directly into a CREATE INDEX SQL statement via string formatting, without identifier quoting or allowlist validation.

What are the systems affected?

The following version(s) are affected:

- R2R through 3.6.6, in the default configuration (authentication not required on this endpoint) - (Affected)

No fixed version has been published at the time of writing – see “Mitigation process” below

What does this mean?

Typical attack flow:

1. **Submit a crafted index name to the vector index creation endpoint**
    - An unauthenticated attacker sends a request to R2R’s POST /v3/indices endpoint, setting the index_name parameter to a value containing a semicolon followed by additional, attacker-chosen SQL statements rather than a normal identifier.
2. **Execute arbitrary SQL as the database superuser**
    - Because the index name is interpolated directly into a CREATE INDEX statement without quoting or validation, and the database driver accepts multiple semicolon-separated statements in a single query, the attacker’s appended SQL executes with the same PostgreSQL superuser privileges as R2R’s own database connection — granting full read, write, and schema-modification access to the underlying database.

Attack vectors:

- A network-based, unauthenticated request against any internet- or network-reachable R2R deployment running in its default configuration, where the vector index creation endpoint does not require authentication.
- No user interaction, no privileges, and no special access conditions are required (CVSS AV:N/AC:L/PR:N/UI:N). CERTVU is not aware of confirmed active exploitation at the time of writing, but a public proof of concept exists and no fixed version currently exists.

Successful exploitation may allow attackers to:

- Read, modify, or delete any data stored in an affected R2R deployment’s database — including ingested documents, embeddings, user accounts, and access-control configuration — without ever holding a legitimate credential.
- Execute arbitrary schema-modification (DDL) statements with PostgreSQL superuser privileges, and, depending on the database server’s configuration, potentially leverage superuser-level SQL functions to achieve operating-system command execution on the underlying database host.

Mitigation process

CERTVU recommends the following:

1. Require authentication on the affected endpoint immediately — no vendor patch is currently available

- Set require_authentication to true in your R2R deployment configuration so that the vector index creation endpoint (and ideally the API as a whole) cannot be reached by an unauthenticated caller, and restrict network access to the R2R API to trusted hosts only (for

example, via a firewall or reverse proxy) as a further layer of defence. This is a workaround, not a fix — the underlying SQL injection flaw remains present in the application code.

2. Treat this as unpatched: R2R's maintainers had not released a fixed version or formal advisory at the time of writing (GitHub issue #2309 remains open)
3. Remove superuser privileges from R2R's database connection
4. Audit the estate for R2R deployments
5. Review database logs for suspicious index-creation activity

Report suspected compromise to CERTVU at incident@cert.gov.vu or on telephone (678) 33380.

Reference

1. https://www.cve.org/CVERecord?id=CVE-2026-82526
2. https://github.com/SciPhi-AI/R2R/issues/2309